Risk Analysis Using a Hybrid Bayesian-Approximate Reasoning Methodology

T. F. Bott, Los Alamos National Laboratory Los Alamos
S. W. Eisenhawer, Los Alamos National Laboratory

Key Words: Accident Analysis, Bayesian Models, Approximate Reasoning

SUMMARY & CONCLUSIONS

Analysts are sometimes asked to make frequency estimates for specific accidents in which the accident frequency is determined primarily by safety controls. Under these conditions, frequency estimates use considerable expert belief in determining how the controls affect the accident frequency. To evaluate and document beliefs about control effectiveness, we have modified a traditional Bayesian approach by using approximate reasoning (AR)[1] to develop prior distributions. Our method produces accident frequency estimates that separately express the probabilistic results produced in Bayesian analysis and possibilistic results that reflect uncertainty about the prior estimates. Based on our experience using traditional methods, we feel that the AR approach better documents beliefs about the effectiveness of controls than if the beliefs are buried in Bayesian prior distributions. We have performed numerous expert elicitations in which probabilistic information was sought from subject matter experts not trained in probability. We find it much easier to elicit the linguistic variables and fuzzy set membership values used in AR than to obtain the probability distributions used in prior distributions directly from these experts because it better captures their beliefs and better expresses their uncertainties.

1. INTRODUCTION

In this paper, we present a method for generating Bayesian prior distributions of a Poisson parameter using approximate reasoning (AR). In this method, the uncertainty introduced in generating the prior distribution is explicitly represented by fuzzy set memberships interpreted as a possibilistic measure of belief.[2] This method was developed specifically for estimating accident frequencies for military weapon systems in which great reliance is placed on controls to reduce the accident frequency from relatively high to acceptable levels. This approach is useful when there is a lack of "hard" data, but there is a wealth of anecdotal or experiential knowledge. Such a situation arises when experience on a specific weapon system is limited, but more general weapon system experience with safety controls is applicable.

This problem can be approached probabilistically using Bayesian statistical analysis.[3] To review this concept briefly, subjective estimates of the Poisson parameter called prior distributions are "updated" using available operating data to produce an updated estimate of the parameter. When there is little operating experience or useful surrogate data, the prior distribution can dominate the results. Such prior distributions often are generated using expert judgement that is difficult to document, and the original justification may be lost.

In the work reported here, we use Bayesian methods to include nonstatistical knowledge about the effect of safety controls on accident frequency. An important innovation is the use of the mathematical tools of AR to capture the knowledge base and reasoning used by experts in constructing prior distributions. This approach provides a rigorous, reproducible, and traceable basis for the prior distributions. It also provides a means for explicitly indicating uncertainty about the prior distribution using possibility as an uncertainty measure. In a typical Bayesian analysis, the uncertainty about the prior distribution is folded into the distribution itself, a practice that tends to obscure the issues involved in generating the prior distribution. In our method, this source of uncertainty is treated separately and differently from probabilistic uncertainty

[1]For a good survey of this field see Ramon Lopez de Mantaras, *Approximate Reasoning Models*, Ellis Horwood Series in Artificial Intelligence, Ellis Howrood LTD, 1990.

[2]D. Dubois and H. Prade, *Possibility Theory*, Plenum Press, 1988.

[3]H. F. Martz and R. A. Waller, *Bayesian Reliability Analysis*, John Wiley and Sons, 1982.

by interpreting fuzzy set memberships a as a measure of the expert's uncertainty in generating prior distributions.

A schematic overview of our approach is shown in Fig. 1. In this paper, we focus mainly on evaluating controls, generating prior distributions, and generating occurrence probability estimates. The logical decomposition of an event into causal sequences is a critical aspect of our analysis because it allows experts to consider individual sequences leading to an accident one at a time. This simplification is necessary in identifying the controls used to prevent an accident and in determining their effectiveness. We do not discuss this important part of the analysis here but refer the reader to other discussion of this subject.[4] We also do not discuss Bayesian analysis in detail because this technique is familiar to practitioners of reliability and probabilistic safety analysis. We will spend the majority of this paper describing the AR evaluation of control effectiveness and translating this evaluation into the λ prior distributions. We also will show how the possibilistic measures of uncertainty introduced by the AR analysis are propagated to the occurrence probability estimates. These measures of uncertainty capture the expert's beliefs about the effectiveness of the controls used to reduce accident frequency.

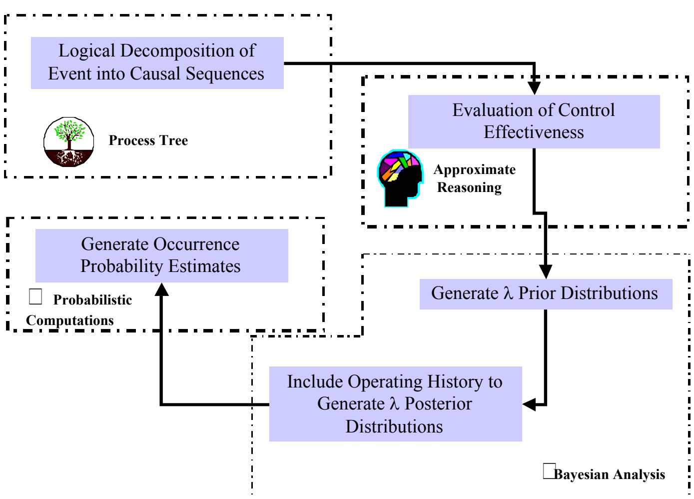

Fig. 1. Overview of the Approach.

[4]S. W. Eisenhawer and T. F. Bott, "Application of Approximate Reasoning to Safety Analysis, "17th *International System Safety Conference*, System Safety Society, August, 1999, Orlando, Florida, Los Alamos National Laboratory report LA-UR-99-1932.

2. ILLUSTRATIVE EXAMPLE DEFINITION

The technique shown in this paper is nearly impossible to follow without an example. The actual problems for which we developed and applied this technique are classified, so we are unable to discuss either the accident sequences or the numerical results in an open forum. However, the example that we use to illustrate our method captures the important characteristics of the actual applications.

We are interested in estimating the probability of occurrence of a particular accident state for a weapon system during a time period τ. We assume that the occurrence can result from any of four independent sequences of events. Each sequence n can be modeled as a Poisson process with a constant occurrence rate λ_n. For simplicity, we assume that no occurrences of the event have happened but presume considerable qualitative knowledge about the controls used to reduce the frequencies of the various sequences leading to the accident conditions.

Controls for the four accident sequences are summarized and evaluated in Table 1. In an actual application, the controls would be identified and evaluated by weapon system experts. These evaluations use an agreed-upon set of linguistic descriptors for control effectiveness, in this case {Highly, Quite, Partially, Ineffective}. This set of descriptors is called a Universe of Discourse (UOD). These descriptors are defined in Table 2.

3. USING THE CONTROLS TO ESTIMATE λ PRIOR DISTRIBUTIONS

The evaluation of the effectiveness of controls forms the basis of our estimation of the λ priors for a Bayesian analysis. We are going to treat these effectiveness descriptors as linguistic variables and fuzzy subsets of the UOD. Our next step is to translate the qualitative descriptions of Table 1 into fuzzy set membership vectors. This translation is called set assignment and introduces set assignment uncertainty. A fuzzy set membership vector shows the set membership values for each of the fuzzy subsets for control effectiveness in the order: {Highly, Quite, Partially, Ineffective}. For example, a vector representing membership of .5 in Highly and .5 in Quite, with no membership in Partially or Ineffective would be {.5,.5,0,0}. The mapping from the qualitative descriptions to the membership vectors relies on the judgment of the analyst, but we have found it helpful to provide the set of guidelines shown in Table 3.

Following these guidelines, the qualitative descriptions of Table 1 translate into the fuzzy set membership vectors shown in Table 4. These fuzzy set

Table 1
Description of Controls for Illustrative Example

Sequence	Controls	Experts Beliefs concerning the Effectiveness of Controls
1	1.1	Highly effective with high confidence
2	2.1	Favor Highly effective but could only be Quite effective
3	3.1	Favor Highly effective but could only be Quite effective
	3.2	No preference for Quite or Partially
4	None	Initiating event is lightning strike which has frequency of about 1×10^{-4} per year

Table 2
Definition of Control Effectiveness Linguistics

Effectiveness Descriptor	Definition
Highly Effective	The control virtually eliminates the occurrence of the sequence
Quite Effective	The control greatly reduces the occurrence rate of the sequence
Partially Effective	The control somewhat reduces the occurrence rate of the sequence
Ineffective	The control does not affect the occurrence rate of the sequence

Table 3
Set Assignment Membership Value Guidelines

Belief Description	Set Assignment Value	Complementary Set Assignment Values
Belief that value is exclusively in one set	1	0
Strong belief that value is in one set, but some belief that another set may also be appropriate	0.9	0.1
Equal belief that the value is in any of n sets	1/n	1/n for each
One set is favored, but another has significant support as well	0.7	0.3

Table 4
Control Effectiveness Fuzzy Set Membership Vectors

Control	Control Effectiveness Linguistic Descriptor			
	Highly	Quite	Partially	Ineffective
1.1	1	0	0	0
2.1	.7	.3	0	0
3.1	.7	.3	0	0
3.2	0	.5	.5	1
Combined 3.1 and 3.2	.5	.3	0	0
Effective 4	0	1	0	0

membership values will be interpreted as expressing the expert's belief in which prior estimates for λ to use for each sequence. The greater the fuzzy set membership, the greater the expert's belief that a given effectiveness is appropriate for a set of controls.

In our example, two controls, 3.1 and 3.2, are used to reduce the accident frequency for sequence 3. To apply our method in such a situation, the analysts evaluate the aggregate effectiveness of the control suite using a rule base.[5] An example of such a rule base is shown in Table 5. This rule base accepts control effectiveness descriptors for two controls and outputs the effectiveness of the combination using the same linguistic descriptors as the input. The rule base shown here is used commonly in our analysis and represents a slight bias toward conservatism in combining controls. This bias is seen in the result for two partially effective controls. Two partially effective controls result in a partially effective aggregate control. This rule prevents stringing together a series of mediocre controls and claiming that the result is Highly effective.

The operation of this rule base is illustrated by evaluating the effectiveness of controls 3.1 and 3.2. Control 3.1 has an effectiveness described by the set membership vector {0.7,0.3,0,0} and 3.2 has {0,0.5,0.5,0}. Recall that the first position in the vector is for Highly, the second is for Quite, the third is for Partially, and the fourth is for Ineffective. The Cartesian product of the membership vectors for controls 3.1 and 3.2 generates all the combinations of control effectiveness that have non-zero memberships in both controls. The pairs of effectiveness descriptors that we have to consider are (Highly, Quite), (Highly, Partially), (Quite, Partially), and (Quite, Quite). Our notation is that the first value is from control 3.1, and the second is from 3.2. As an example, according to the rule base Highly and Quite effective controls combine to produce a Highly effective aggregate. The inferences of interest in our example are shown by highlighting the appropriate items in the rule base. Note that three of the pairs result in an output of Highly.

The control effectiveness descriptors are fuzzy sets and have memberships associated with them. We need some way to generate the membership value of the output effectiveness descriptors from the membership values of the inputs. The membership value for the resultant arising from a pair of inputs is found by taking the minimum of the memberships values for the pair. This works fine when there is only one pair of inputs that results in a given output. However, in our case, there are three pairs of input that lead to the same output, namely, Highly. In this case, the membership of the output is found using the Max-Min formula.[6] This formula is succinctly stated as

$$\mu_{\Re} = \underset{\forall (n,m) \to \Re}{Max} \left(Min\left(\kappa_n, \sigma_m\right)\right) \quad (1)$$

In this formula, κ_n and σ_m are elements n and m of fuzzy input membership vectors κ and σ and $\Re$ is a particular element output by the rule. To find the membership for $\Re$, one first finds the minimum membership in either κ or σ for every pair that result in $\Re$. The membership value of the resultant $\Re$ is then the maximum value over all pairs of the inputs κ and σ that result in $\Re$.

An example using the rule base of Table 5 and the membership values for controls 3.1 and 3.2 is shown in Table 6. Using the Max-Min formula, the resultant membership vector has values {.5,.3,0,0}, indicating a stronger belief in Highly than in Quite and no belief in Partially or Ineffective.

The final complication in our example is sequence 4. In this sequence, there are no controls, but the sequence frequency has a relatively low inherent frequency. In some sequences, constraints or other factors not normally considered controls may dictate the frequency of the sequence. We often encounter sequences whose frequencies are dictated primarily by the occurrence of external initiating events such as lightning. The effect of the relative rarity of lightning strikes on the system may be treated as if it were a control, and can even be combined with other controls. As we shall demonstrate later, the sequence 4 inherent frequency of about 10^{-4} per year corresponds to a control with that has full membership in Quite effective.

4. BAYES PRIOR ESTIMATE FOR POISSON PARAMETERS

As stated above, we have assumed that the occurrence of each sequence can be described using a Poisson process with the occurrence times distributed exponentially according to a Poisson parameter λ_n for sequence n. To use a Bayesian estimation process, we make an initial or prior estimate for each λ using existing knowledge about each sequence and then modify that prior estimate using occurrence data derived from operational experience. In this example, we assume that

[5]T. F. Bott, "An Approach to Evaluating the Effectiveness of Safety Controls," Los Alamos National Laboratory report LA-UR-98-4953 (1998).

[6]T. J. Ross, *Fuzzy Logic with Engineering Applications*, McGraw-Hill, New York, 1995.

Table 5
Rule Base for Combining Reinforcing Controls

Control 3.2 Effectiveness		**Control 3.1 Effectiveness**			
		Ineffective	*Partially*	*Quite*	*Highly*
	Ineffective	Ineffective	Partially	Quite	Highly
	Partially	Partially	Partially	Quite	Highly
	Quite	Quite	Quite	Highly	Highly
	Highly	Highly	Highly	Highly	Highly

Table 6
Effectiveness Membership Values for Combined Controls

Control 3.2 Effectiveness		**Control 3.1 Effectiveness**			
		Ineffective (0)	*Partially* (0)	*Quite* (.3)	*Highly* (.7)
	Ineffective (0)	Ineffective	Partially	Quite	Highly
	Partially (.5)	Partially	Partially	Quite Min(.3,.5)→ .3	Highly Min(.7,.5)→ .5
	Quite (.5)	Quite	Quite	Highly Min(.3,.5)→ .3	Highly Min(.7,.5)→ .5
	Highly (0)	Highly	Highly	Highly	Highly

we can assign λ to intervals. The prior estimates are assumed to be uniform distributions over these intervals. Although the assumption of a uniform distribution is not necessary to use this AR approach, we feel that such a choice generally will be appropriate to the level of knowledge we are assuming in using this method. If enough knowledge exists to make more detailed prior estimates, then the AR approach probably does not use all the available information efficiently.

Using a uniform prior distribution for λ on the interval $[\lambda_2,\lambda_1]$ and no occurrences during a time τ, a Bayes formula produces a posterior distribution for λ given by

$$g(\lambda \mid 0) = \frac{\tau e^{-\lambda\tau}}{e^{-\lambda_2 \tau} - e^{-\lambda_1 \tau}} \quad .(2)$$

Statistics for λ can be generated from this distribution. For example, the mean λ for a given distribution g is given by

$$\bar{\lambda} = \frac{1}{\tau}\left[1 + \frac{\lambda_2 \tau e^{-\lambda_2 \tau} - \lambda_1 \tau e^{-\lambda_1 \tau}}{e^{-\lambda_2 \tau} - e^{-\lambda_1 \tau}}\right] \quad .(3)$$

Similarly, the γ probability value for λ is given by

$$\lambda_\gamma = -\frac{\ln\left[(1-\gamma)e^{-\lambda_2 \tau} - \gamma e^{-\lambda_1 \tau}\right]}{\tau} \quad .(4)$$

Although Eq. (4) formally represents a probability interval for λ, we will refer to it as the γ^{th} percentile. These formulae are applicable to each sequence leading to the accident conditions. The formulae will depend only on which λ interval is chosen as a prior.

5. BAYES PRIOR ESTIMATES VIEWED AS FUZZY SET ASSIGNMENTS

To translate control effectiveness into λ intervals, we define a set of intervals on the real line that correspond to the definitions for each effectiveness descriptor in Table 2. For simplicity, we wish to have a one-to-one mapping between λ intervals and effectiveness descriptors. Thus, each effectiveness will map into a single λ interval. We typically base our definitions on the probability of one or more accident events occurring during a time period t given a particular λ. This probability is found from

$$P[s \geq 1 \; in \; t] = 1 - e^{-\lambda t} \quad .(5)$$

We use the system design lifetime for the value t—in our example about 20 yr. In the systems we have examined, most accident sequences would be expected to have a relatively high probability of occurrence during the design lifetime in the absence of controls. To capture this, we define the lower bound of our highest λestimate interval, I_4, at about .03 so that the probability of occurrence in the design life is $P \approx 0.5$. This definition also fixes the upper bound of our second highest interval, I_3. A highly effective control "virtually eliminates" an accident sequence. We consider an accident sequence as being virtually eliminated if the probability of occurrence during design lifetime is less than about 10^{-3}. This sets the upper bound of our lowest interval I_1 at about 3×10^{-5}/yr. We choose the upper limits on the remaining interval, I_2, corresponding to a quite effective control as .003, 2 orders of magnitude higher than the I_1 upper bound and an order of magnitude below the I_3 upper bound. The results of the analysis are insensitive to the lower limit on the intervals I_1 to I_3 as long as the interval covers a decade or more. Therefore, we typically choose the lower bound for the lowest interval as 10^{-6}. This set of intervals and the mapping from the effectiveness descriptors to the intervals is shown in Fig. 2.

This mapping emulates how an expert perceives control effectiveness affecting an initial estimate of λ in the interval I_4. A control that is ineffective leaves λ in the I_4 interval. A partially effective control moves λ to I_3, Quite effective moves it to I_2 and Highly effective moves it to I_1.

The one-to-one mapping between control effectiveness and λ intervals means that the control effectiveness fuzzy membership values can be assigned one-to-one to λ intervals. We interpret these fuzzy set membership assignments for a λ interval as indicating

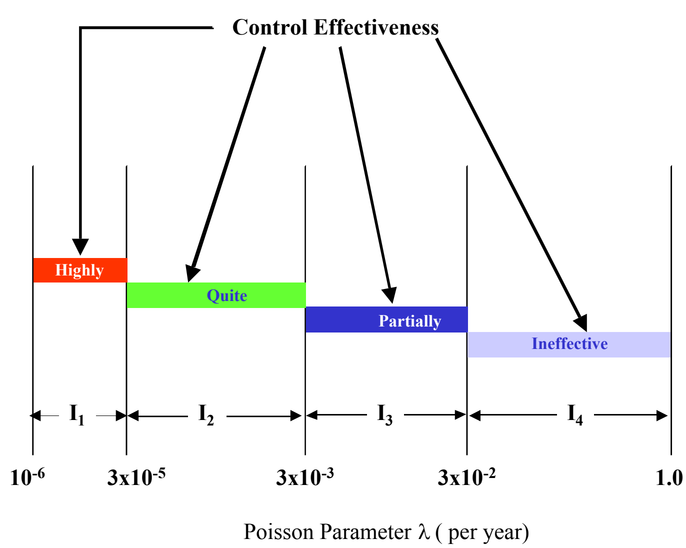

Fig. 2. Mapping from Control Effectiveness Fuzzy Subsets to λ Intervals.

our belief that a given λ interval is the appropriate one to represent the accident-sequence frequency, when the controls are taken into account.

Recall that each of the intervals I_1 through I_4 is a uniform distribution for λ and can be used in Eq. (2) to generate a posterior λ distribution. Each of these λ distributions can be used to generate λ statistics, which we then use to generate probabilities of occurrence using Eq. (5). In Table 7, we summarize the mean λ, 90^{th} percentile λ, and the probabilities of occurrence for a 20-yr design lifetime that are generated by the different λ intervals.

To illustrate, consider sequence 2 of our example. The control effectiveness membership vector is {.7,.3,0,0}. We interpret this to mean that we have a relative belief value .7 that λ_4 is in I_1, .3 that it is in I_2, and 0 that it is in I_2 or I_3. This means that we have the relative belief that mean λ_4 is 1.5×10^5 and $P_{s\geq1}$ is 3×10^{-4}. Similarly, we have a belief value of 0.3 that the mean λ_4 is 1.5×10^3 and $P_{s\geq1}$ is 3×10^{-2}.

The result then carries both the uncertainty associated with the λ prior estimate (represented by a uniform distribution over an interval) and uncertainty on assigning a λ prior interval by means of control effectiveness evaluation. The former uncertainty is expressed by the $g(\lambda|0)$ distribution, and the latter is expressed by the fuzzy subset memberships (interpreted as belief) associated with the $g(\lambda|0)$ distribution.

6. INTERPRETING THE PROBABILITY OF OCCURRENCE RESULTS

Some representative probability of occurrence results for our example are shown in Table 8. The upper bound uses the mean λ from the posterior distribution generated by the highest λ interval with non-zero set membership. The lower bound is the result of using the λ posterior distribution generated by the lowest λ interval with non-zero set membership. When only an upper bound is given, only one of the λ intervals has non-zero set membership. The best estimate is found by:

❶ If one of the λ set memberships is maximal use that λ, or

❷ If the two largest λ set memberships are tied (typically both at .5) then use the geometric mean of the λ's given by

$$\bar{\lambda}_{1,2} = (\lambda_1 \lambda_2)^{1/2} \qquad (7)$$

In Table 5, sequence 2 exercises rule ❶ and sequence 3 with control 3.2 only exercises rule ❷.

Table 7
Mean and 90th Percentile Values for λ

Prior λ Interval	Mean		90th Percentile	
	λ	$P_{s\geq1}$	λ	$P_{s\geq1}$
I_1 [1×10^{-6}, 3×10^{-5}]	1.5×10^{-5}	3×10^{-4}	2.7×10^{-5}	5.4×10^{-4}
I_2 [1×10^{-6}, 3×10^{-5}]	1.5×10^{-3}	3×10^{-2}	2.7×10^{-3}	5×10^{-2}
I_3 [1×10^{-6}, 3×10^{-52}]	1.4×10^{-2}	2.4×10^{-1}	2.6×10^{-2}	4×10^{-1}
I_4 [3×10^{-62}, 1]	1.0×10^{-1}	8.6×10^{-1}	1.9×10^{-1}	9.8×10^{-1}

Table 8
Representative Results for Individual Sequence Probabilities of Occurrence

Possible Event Sequences	Mean λ Occurrence Probability Interval		
	Lower Bound	Upper Bound	Best Estimate
Sequence 1	-	.0003	.0003
Sequence 2	.0003	.03	.0003
Sequence 3 Control 3.2 only	03	.24	.18

The probability of one or more occurrences from any of n sequences is found from the formula

$$P[s \geq 1 \; in \; t] = 1 - e^{-\sum_n \lambda_n t} \quad (5)$$

This is easily and rapidly computed using a Monte Carlo or other sampling simulation.

One potential drawback to this approach is the added complexity of the results. Our accident frequency estimates include both the Bayesian distribution and fuzzy set memberships interpreted as possibilities or beliefs. Potential users of safety results often wish to get a single, bottom-line answer, not a proliferation of uncertainty measures. We have addressed this by explaining the interaction of the beliefs and the statistics. For example, we describe the results of Table 8 as showing that we are quite certain that the probability of occurrence for sequence 2 is less than .03 and that the average is .0003 or less. This has mollified our sponsors to some extent, but we feel that better methods of communicating the results are needed.

BIOGRAPHIES

Stephen W. Eisenhawer, Ph.D., P.E.
MS K557
Los Alamos National Laboratory
Los Alamos, NM 87545
Internet(e-mail): seisenhawer@lanl.gov

Dr. Eisenhawer's research interests have included multi-phase flow and enhanced oil recovery methods. His current research is in the use of formal logic models and imprecise uncertainty measures for decision analysis.

Terry F. Bott, Ph.D.
MS K557
Los Alamos National Laboratory
Los Alamos, NM 87545
Internet(e-mail): tbott@lanl.gov

Dr. Bott has worked in safety research and applications for 25 years. His current interests include applications of logic models and approximate reasoning to decision making in complex systems

Risk Analysis Using a Hybrid Bayesian-Approximate Reasoning Method

T. F. Bott and S. W. Eisenhawer

Decision Applications Division

Los Alamos National Laboratory

Problem Definition

Estimate Probability of Accident Occurrence

Problem Attributes

- Accident can result from several causal sequences
- Frequency is dominated by Safety Controls
- Event can be modeled as a Poisson Process
- Limited operating experience and surrogate data sources

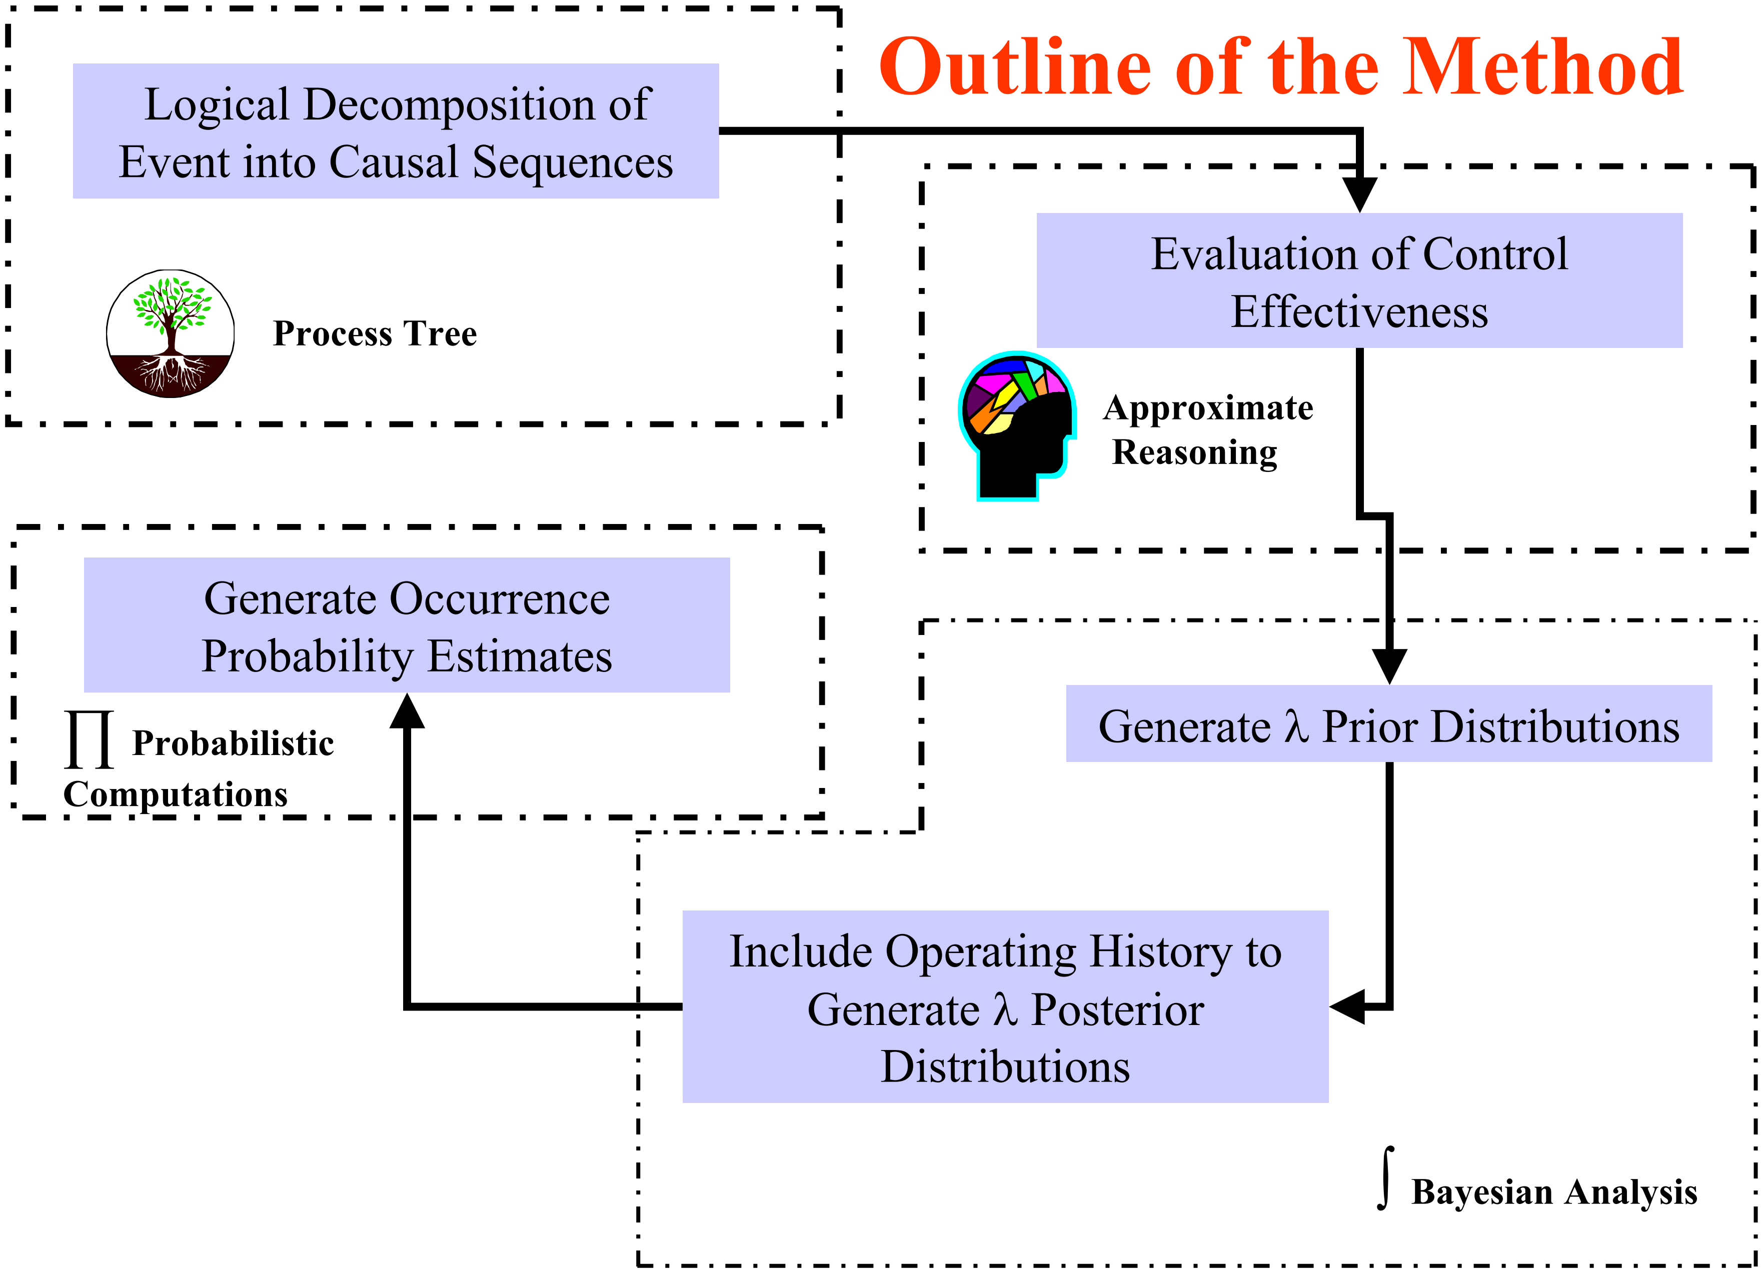
Outline of the Method
Logical Decomposition of Event into Causal Sequences
Process Tree
Evaluation of Control Effectiveness
Approximate Reasoning
Generate Occurrence Probability Estimates
∏ Probabilistic Computations
Generate λ Prior Distributions
Include Operating History to Generate λ Posterior Distributions
∫ Bayesian Analysis

Evaluation of Control Effectiveness

Definition of Control Effectiveness Linguistics

Effectiveness Descriptor	Definition
Highly Effective	The control virtually eliminates the occurrence of the sequence
Quite Effective	The control greatly reduces the occurrence rate of the sequence
Partially Effective	The control somewhat reduces the occurrence rate of the sequence
Ineffective	The control does not affect the occurrence rate of the sequence

Description of Controls for Illustrative Example

Sequence	Controls	Experts Beliefs concerning the Effectiveness of Controls
1	1.1	Highly effective with high confidence
2	2.1	Favor Highly effective but could only be Quite effective
3	3.1	Favor Highly effective but could only be Quite effective
	3.2	No preference for Quite or Partially
4	None	Initiating event is lightning strike which has frequency of about 1×10^{-4} per year

Evaluation of Control Effectiveness

Set Assignment Membership Value Guidelines

Belief Description	Set Assignment Value	Complementary Set Assignment Values
Belief that value is exclusively in one set	1	0
Strong belief that value is in one set, but some belief that another set may also be appropriate	0.9	0.1
Equal belief that the value is in any of n sets	1/n	1/n for each
One set is favored, but another has significant support as well	0.7	0.3

Control Effectiveness Fuzzy Set Membership Vectors

Control	Control Effectiveness Linguistic Descriptor			
	Highly	Quite	Partially	Ineffective
1.1	1	0	0	0
2.1	.7	.3	0	0
3.1	.7	.3	0	0
3.2	0	.5	.5	1
Combined 3.1 and 3.2	.5	.3	0	0
Effective 4	0	1	0	0

Evaluation of Control Effectiveness

Effectiveness Membership Values for Combined Controls

	Control 3.1 Effectiveness			
	Ineffective (0)	*Partially* (0)	*Quite* (.3)	*Highly* (.7)
Ineffective (0)	Ineffective	Partially	Quite	Highly
Partially (.5)	Partially	Partially	Quite Min(.3,.5)→ .3	Highly Min(.7,.5)→ .5
Quite (.5)	Quite	Quite	Highly Min(.3,.5)→ .3	Highly Min(.7,.5)→ .5
Highly (0)	Highly	Highly	Highly	Highly

Fuzzy Set Representation of Control Effectiveness

Generate λ Prior Distributions

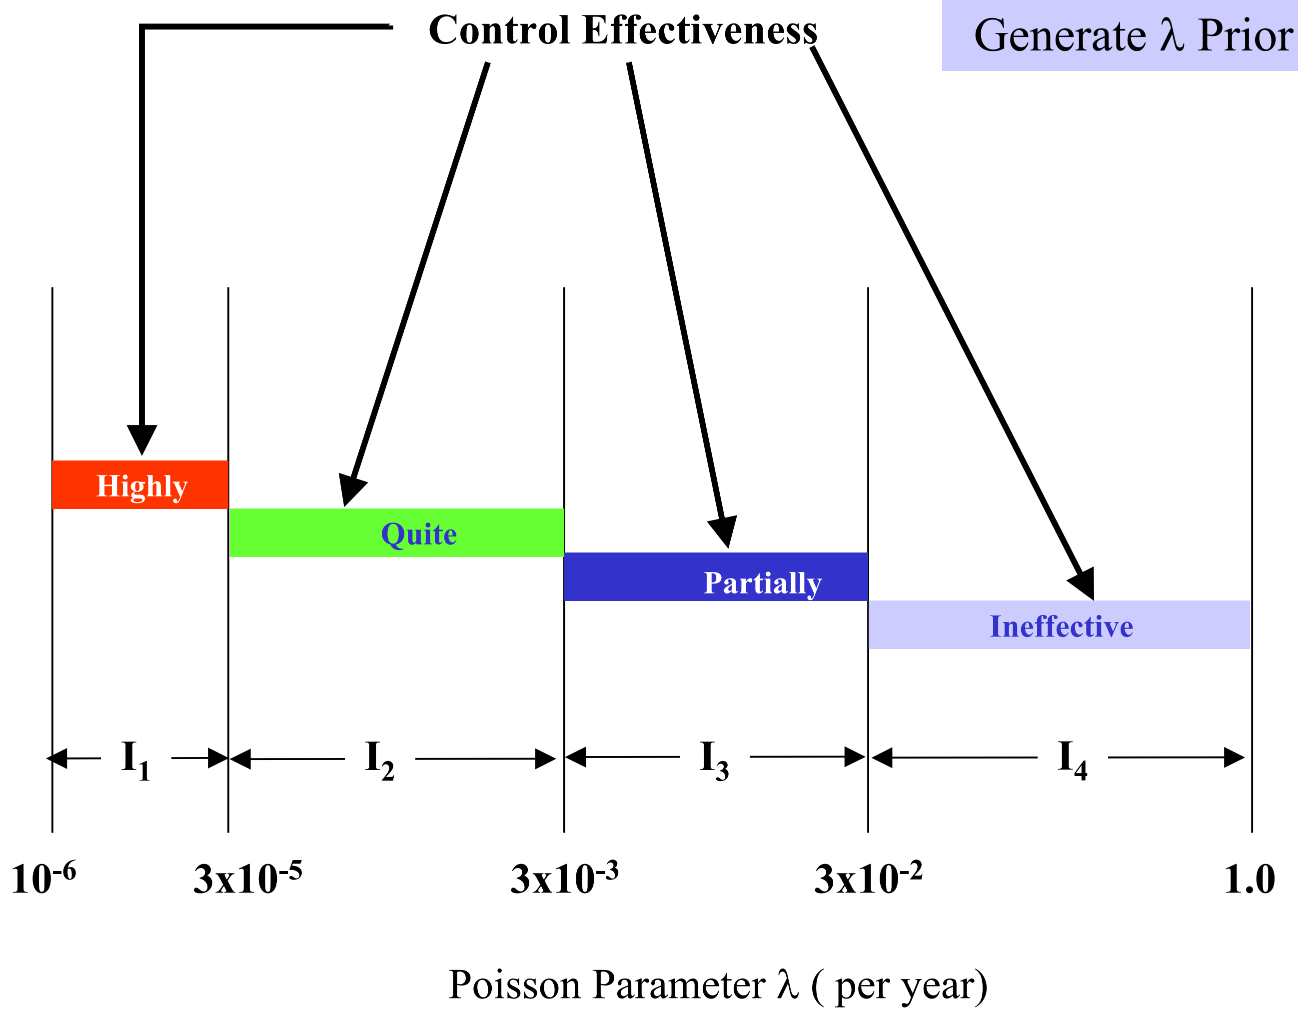

Generate λ Prior Distributions

$$g(\lambda \mid 0) = \frac{\tau e^{-\lambda\tau}}{e^{-\lambda_2\tau} - e^{-\lambda_1\tau}}$$

λ Prior Distributions for Poisson Parameter with Uniform Prior

Mean λ Generated by Prior and Operating Experience

$$\bar{\lambda} = \frac{1}{\tau}\left[1 + \frac{\lambda_2 \tau e^{-\lambda_2\tau} - \lambda_1 \tau e^{-\lambda_1\tau}}{e^{-\lambda_2\tau} - e^{-\lambda_1\tau}}\right]$$

$$\lambda_\gamma = -\frac{\ln\left[(1-\gamma)e^{-\lambda_2\tau} - \gamma e^{-\lambda_1\tau}\right]}{\tau}$$

γth Cumulative Probability
λ Generated by Prior and Operating Experience

Generate Occurrence Probability Estimates

Probability of One or More Occurrences in Time t

$$P[s \geq 1 \ in \ t] = 1 - e^{-\lambda t}$$

Generate Occurrence Probability Estimates

$$\bar{\lambda} = \frac{1}{\tau}\left[1 + \frac{\lambda_2 \tau e^{-\lambda_2 \tau} - \lambda_1 \tau e^{-\lambda_1 \tau}}{e^{-\lambda_2 \tau} - e^{-\lambda_1 \tau}}\right]$$

$$P[s \geq 1 \ in \ t] = 1 - e^{-\lambda t}$$

Mean and 90th Percentile Values for λ

Prior λ Interval	**Mean**		**90th Percentile**	
	λ	**$P_{s\geq1}$**	**λ**	**$P_{s\geq1}$**
I_1 [1 x 10^{-6}, 3 x 10^{-5}]	1.5 x 10^{-5}	3 x 10^{-4}	2.7 x 10^{-5}	5.4 x 10^{-4}
I_2 [1 x 10^{-6}, 3 x 10^{-5}]	1.5 x 10^{-3}	3 x 10^{-2}	2.7 x 10^{-3}	5 x 10^{-2}
I_3 [1 x 10^{-6}, 3 x 10^{-52}]	1.4 x 10^{-2}	2.4 x 10^{-1}	2.6 x 10^{-2}	4 x 10^{-1}
I_4 [3 x 10^{-62}, 1]	1.0 x 10^{-1}	8.6 x 10^{-1}	1.9 x 10^{-1}	9.8 x 10^{-1}

Include Operating History to Generate λ Posterior Distribution Statistics

$$\lambda_\gamma = -\frac{\ln\left[(1-\gamma)e^{-\lambda_2 \tau} - \gamma e^{-\lambda_1 \tau}\right]}{\tau}$$

Generate Occurrence Probability Estimates

Possible Event Sequences	Mean λ Occurrence Probability Interval		
	Lower Bound	Upper Bound	Best Estimate
Sequence 1	-	.0003	.0003
Sequence 2	.0003	.03	.0003
Sequence 3 Control 3.2 only	03	.24	.18

Multiple Probability Estimates Reflect Uncertainty Generated during Construction of Prior Distribution

Evaluation of Control Effectiveness

Quite and Partially Produces Quite

		Control 3.1 Effectiveness			
		Ineffective (0)	*Partially* (0)	*Quite* (.3)	*Highly* (.7)
	Ineffective *(0)*	Ineffective	Partially	Quite	Highly
	Partially (.5)	Partially	Partially	Quite Min(.3,.5)$\rightarrow$ .3	Highly Min(.7,.5)$\rightarrow$ .5
	Quite (.5)	Quite	Quite	Highly Min(.3,.5)$\rightarrow$ .3	Highly Min(.7,.5)$\rightarrow$ .5
	Highly (0)	Highly	Highly	Highly	Highly

Rule Base for Evaluating the Effectiveness of Multiple Controls

Maximum of Minima for Highly Produces Membership of .5 in Highly

Summary and Conclusions

Provides a structured Method for Constructing Prior Distributions

Provides a Traceable Documentation Trail for Prior Distributions

Provides Separate Uncertainty Measure for Prior Distribution

Efficient Method for Collecting Expert Judgment on Control Effectiveness